

CLAIMS

1. A method of processing data in a computer system comprising at least one host and at least one content addressable storage system which stores data for the at least one host, wherein the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, the method comprising an act of:
 - (a) in response to an access request from the at least one host computer for a unit of data identified by a content address, parsing the content address to determine at least one aspect of a physical storage location for the unit of data on the at least one storage system.
2. The method of claim 1, wherein the at least one storage system includes a plurality of storage nodes, and wherein the act (a) further comprises an act of parsing the content address to determine which of the plurality of storage nodes includes the physical storage location for the unit of data.
3. The method of claim 2, wherein at least some of the plurality of storage nodes include a plurality of disks, and wherein the act (a) further comprises an act of parsing the content address to determine which of the plurality of disks includes the physical storage location for the unit of data.
4. The method of claim 1, wherein the act (a) is performed in response to a request to retrieve the unit of data from the at least one storage system, and wherein the method further comprises an act of passing the unit of data to the at least one host.
5. The method of claim 1, wherein the act (a) is performed in response to a request to write the unit of data to the at least one storage system.
6. The method of claim 5, further comprising an act of storing the unit of data at least partially at the physical storage location.
7. The method of claim 5, further comprising acts of:

applying an algorithm to determine a specified physical storage location based on the content address;

determining whether the specified physical storage location is suitable to store the unit of data, and when it is not, performing acts of:

storing the unit of data at a different physical storage location; and

storing a pointer to the different physical storage location at the specified physical storage location.

8. The method of claim 7, further comprising acts of:

moving the unit of data from the different physical storage location to the specified storage location; and

deleting the pointer to the different physical storage location.

9. The method of claim 1, wherein the storage system comprises a plurality of storage nodes, and wherein the method further comprises an act of assigning, to at least one of the plurality of storage nodes, a range of content addresses so that the at least one of the plurality of storage nodes is assigned to store a plurality of units of data having content address within the range of content addresses.

10. The method of claim 1, further comprising an act of determining the physical storage location of the unit of data solely by the act of parsing and without performing an index lookup.

11. At least one computer readable medium encoded with instructions that, when executed on a computer system perform, a method of processing data, wherein the computer system comprises at least one host and at least one content addressable storage system which stores data for the at least one host, and wherein the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, the method comprising an act of:

(a) in response to an access request from the at least one host computer for a unit of data identified by a content address, parsing the content address to determine at least

one aspect of a physical storage location for the unit of data on the at least one storage system.

12. The at least one computer readable medium of claim 11, wherein the at least one storage system includes a plurality of storage nodes, and wherein the act (a) further comprises an act of parsing the content address to determine which of the plurality of storage nodes includes the physical storage location for the unit of data.

13. The at least one computer readable medium of claim 12, wherein at least some of the plurality of storage nodes include a plurality of disks, and wherein the act (a) further comprises an act of parsing the content address to determine which of the plurality of disks includes the physical storage location for the unit of data.

14. The at least one computer readable medium of claim 11, wherein the act (a) is performed in response to a request to retrieve the unit of data from the at least one storage system, and wherein the method further comprises an act of passing the unit of data to the at least one host.

15. The at least one computer readable medium of claim 11, wherein the act (a) is performed in response to a request to write the unit of data to the at least one storage system.

16. The at least one computer readable medium of claim 15, wherein the method further comprises an act of storing the unit of data at least partially at the physical storage location.

17. The at least one computer readable medium of claim 15, wherein the method further comprises acts of:

applying an algorithm to determine a specified physical storage location based on the content address;

determining whether the specified physical storage location is suitable to store the unit of data, and when it is not, performing acts of:

storing the unit of data at a different physical storage location; and
storing a pointer to the different physical storage location at the specified physical storage location.

18. The at least one computer readable medium of claim 17, wherein the method further comprises acts of:

moving the unit of data from the different physical storage location to the specified storage location; and

deleting the pointer to the different physical storage location.

19. The at least one computer readable medium of claim 11, wherein the storage system comprises a plurality of storage nodes, and wherein the method further comprises an act of assigning, to at least one of the plurality of storage nodes, a range of content addresses so that the at least one of the plurality of storage nodes is assigned to store a plurality of units of data having content address within the range of content addresses.

20. The at least one computer readable medium of claim 11, wherein the method further comprises an act of determining the physical storage location of the unit of data solely by the act of parsing and without performing an index lookup.

21. A content addressable storage system for use in a computer system, including the content addressable storage system and at least one host, wherein the at least one host accesses data units stored on the content addressable storage system using content addresses generated based on the content of the data units, the content addressable storage system comprising:

at least one storage device to store data received from the at least one host; and

at least one controller that, in response to an access request from the at least one host computer for a unit of data identified by a content address, parses the content address to determine at least one aspect of a physical storage location for the unit of data on the at least one storage system.

22. The content addressable storage system of claim 21, further comprising a plurality of storage nodes that comprise the at least one storage device, and wherein the at least one controller parses the content address to determine which of the plurality of storage nodes includes the physical storage location for the unit of data.

23. The content addressable storage system of claim 22, wherein at least some of the plurality of storage nodes include a plurality of disks, and wherein the at least one controller parses the content address to determine which of the plurality of disks includes the physical storage location for the unit of data.

24. The content addressable storage system of claim 21, wherein the at least one controller parses the content address in response to a request to retrieve the unit of data from the at least one storage system, and wherein the controller passes the unit of data to the at least one host.

25. The content addressable storage system of claim 21, wherein the at least one controller parses the content address in response to a request to write the unit of data to the at least one storage system.

26. The content addressable storage system of claim 25, wherein the at least one controller stores the unit of data at the physical storage location.

27. The content addressable storage system of claim 25, wherein the at least one controller:

- applies an algorithm to determine a specified physical storage location based on the content address;

- determines whether the specified physical storage location is suitable to store the unit of data, and when it is not:

- stores the unit of data at a different physical storage location; and

- stores a pointer to the different physical storage location at the specified physical storage location.

28. The content addressable storage system of claim 27, wherein the at least one controller:

moves the unit of data from the different physical storage location to the specified storage location; and

deletes the pointer to the different physical storage location.

29. The content addressable storage system of claim 21, further comprising a plurality of storage nodes that comprise the at least one storage device, wherein the controller assigns, to at least one of the plurality of storage nodes, a range of content addresses so that the at least one of the plurality of storage nodes is assigned to store a plurality of units of data having content address within the range of content addresses.

30. The content addressable storage system of claim 21, wherein the controller determines the physical storage location of the unit of data solely by parsing the content address and without performing an index lookup.

31. A method of processing data in a computer system comprising at least one host and at least one content addressable storage system which stores data for the at least one host, wherein the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, the method comprising acts of:

(a) receiving, from the host, a request to store a unit of data on the storage system, the unit of data having a content address based on the content of the unit of data;

(b) determining, based on the content address, a first storage location on the storage system to which the content address maps;

(c) storing a pointer for the first unit of data at the first storage location, the pointer pointing to a second storage location; and

(d) storing the unit of data at the second storage location on the storage system.

32. The method of claim 31, wherein the act (d) is performed before the acts (b) and (c).

33. The method of claim 31, further comprising acts of:
- (e) receiving, from the host, a request to retrieve the unit of data, the request including a content address of the unit of data;
 - (f) mapping the content address to the first storage location;
 - (g) retrieving the pointer from the first storage location; and
 - (h) using the pointer to access the second storage location and retrieve the unit of data from the second storage location.
34. The method of claim 31, further comprising acts of:
- (i) periodically searching the at least one storage system for pointers to other storage locations on the storage system which store units of data; and
 - (j) determining whether any of the pointers to other storage locations can be replaced with their corresponding units of data.
35. At least one computer readable medium encoded with instructions that, when executed on a computer system, perform a method of processing data, wherein the computer system comprises at least one host and at least one content addressable storage system which stores data for the at least one host, and wherein the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, the method comprising acts of:
- (a) receiving, from the host, a request to store a unit of data on the storage system, the unit of data having a content address based on the content of the unit of data;
 - (b) determining, based on the content address, a first storage location on the storage system to which the content address maps;
 - (c) storing a pointer for the first unit of data at the first storage location, the pointer pointing to a second storage location; and
 - (d) storing the unit of data at the second storage location on the storage system.
36. The at least one computer readable medium of claim 35, wherein the act (d) is performed before the acts (b) and (c).

37. The at least one computer readable medium of claim 35, wherein the method further comprises acts of:

- (e) receiving, from the host, a request to retrieve the unit of data, the request including a content address of the unit of data;
- (f) mapping the content address to the first storage location;
- (g) retrieving the pointer from the first storage location; and
- (h) using the pointer to access the second storage location and retrieve the unit of data from the second storage location.

38. The at least one computer readable medium of claim 35, wherein the method further comprises acts of:

- (i) periodically searching the at least one storage system for pointers to other storage locations on the storage system which store units of data; and
- (j) determining whether any of the pointers to other storage locations can be replaced with their corresponding units of data.

39. A content addressable storage system for use in a computer system that includes at least one host, wherein the at least one host accesses data units stored on the content addressable storage system using content addresses generated based on the content of the data units, the content addressable storage system comprising:

- at least one storage device to store data received from the at least one host; and
- at least one controller that:
 - receives, from the host, a request to store a unit of data on the storage system, the unit of data having a content address based on the content of the unit of data;
 - determines, based on the content address, a first storage location on the storage system to which the content address maps;
 - stores a pointer for the first unit of data at the first storage location, the pointer pointing to a second storage location; and
 - stores the unit of data at the second storage location on the storage system.

40. The content addressable storage system of claim 39, wherein the controller stores the unit of data at the second storage location on the storage system before determining the first storage location and storing the pointer.

41. The content addressable storage system of claim 39, wherein the controller further:

- receives, from the host, a request to retrieve the unit of data, the request including a content address of the unit of data;

- maps the content address to the first storage location;

- retrieves the pointer from the first storage location; and

- uses the pointer to access the second storage location and retrieve the unit of data from the second storage location.

42. The content addressable storage system of claim 39, wherein the controller is adapted to:

- periodically search the at least one storage system for pointers to other storage locations on the storage system which store units of data; and

- determine whether any of the pointers to other storage locations can be replaced with their corresponding units of data.